


JINR cloud service

N. Balashov, A. Baranov, N. Kutovskiy, R. Semenov
Laboratory of Information Technologies, JINR, Dubna, Russia

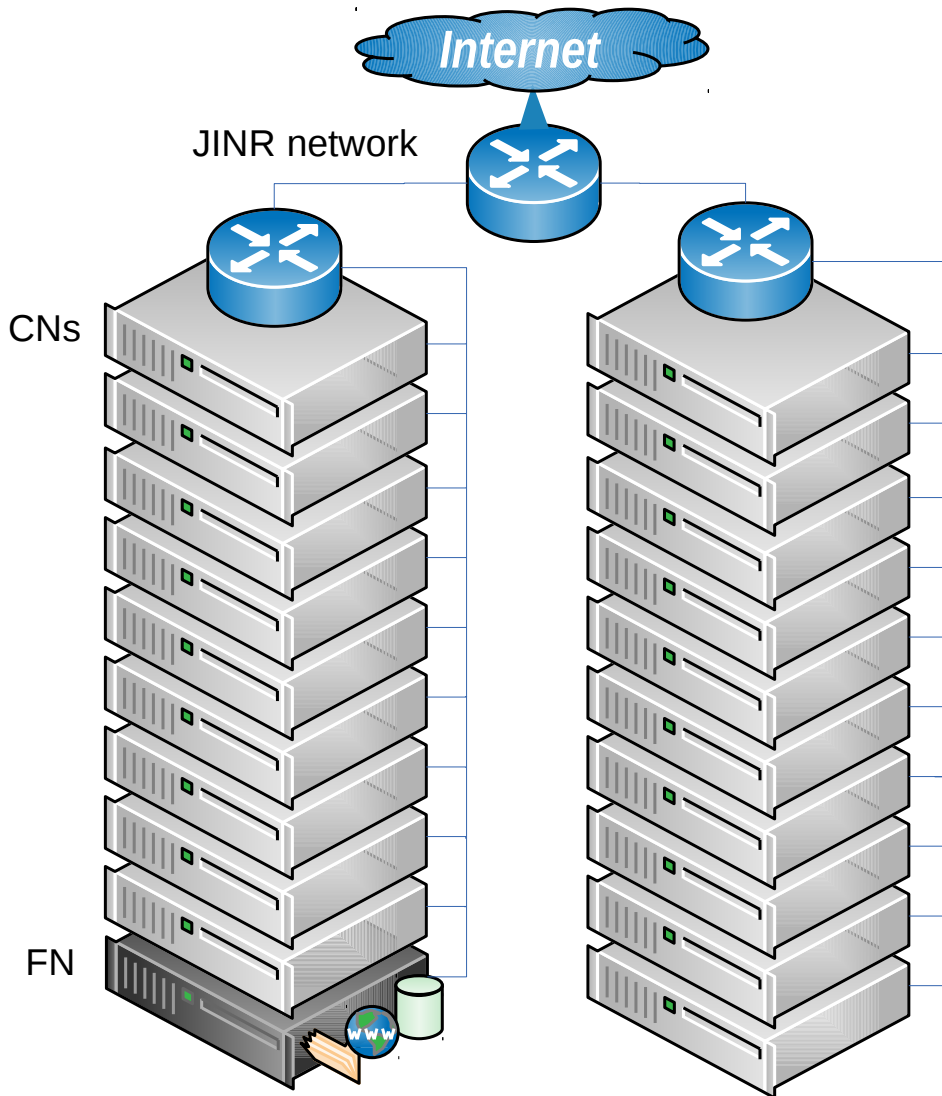
Motivation

- To satisfy the needs of
 - developers (development, testing and debugging various apps in various environments),
 - sysadmins (testing and studying specifics of installation and operation of new apps or testing updates),
 - users (installing and using apps for their daily work)
- To increase an efficiency of hardware utilization
- To simplify access to application software

JINR cloud structure



Implementation



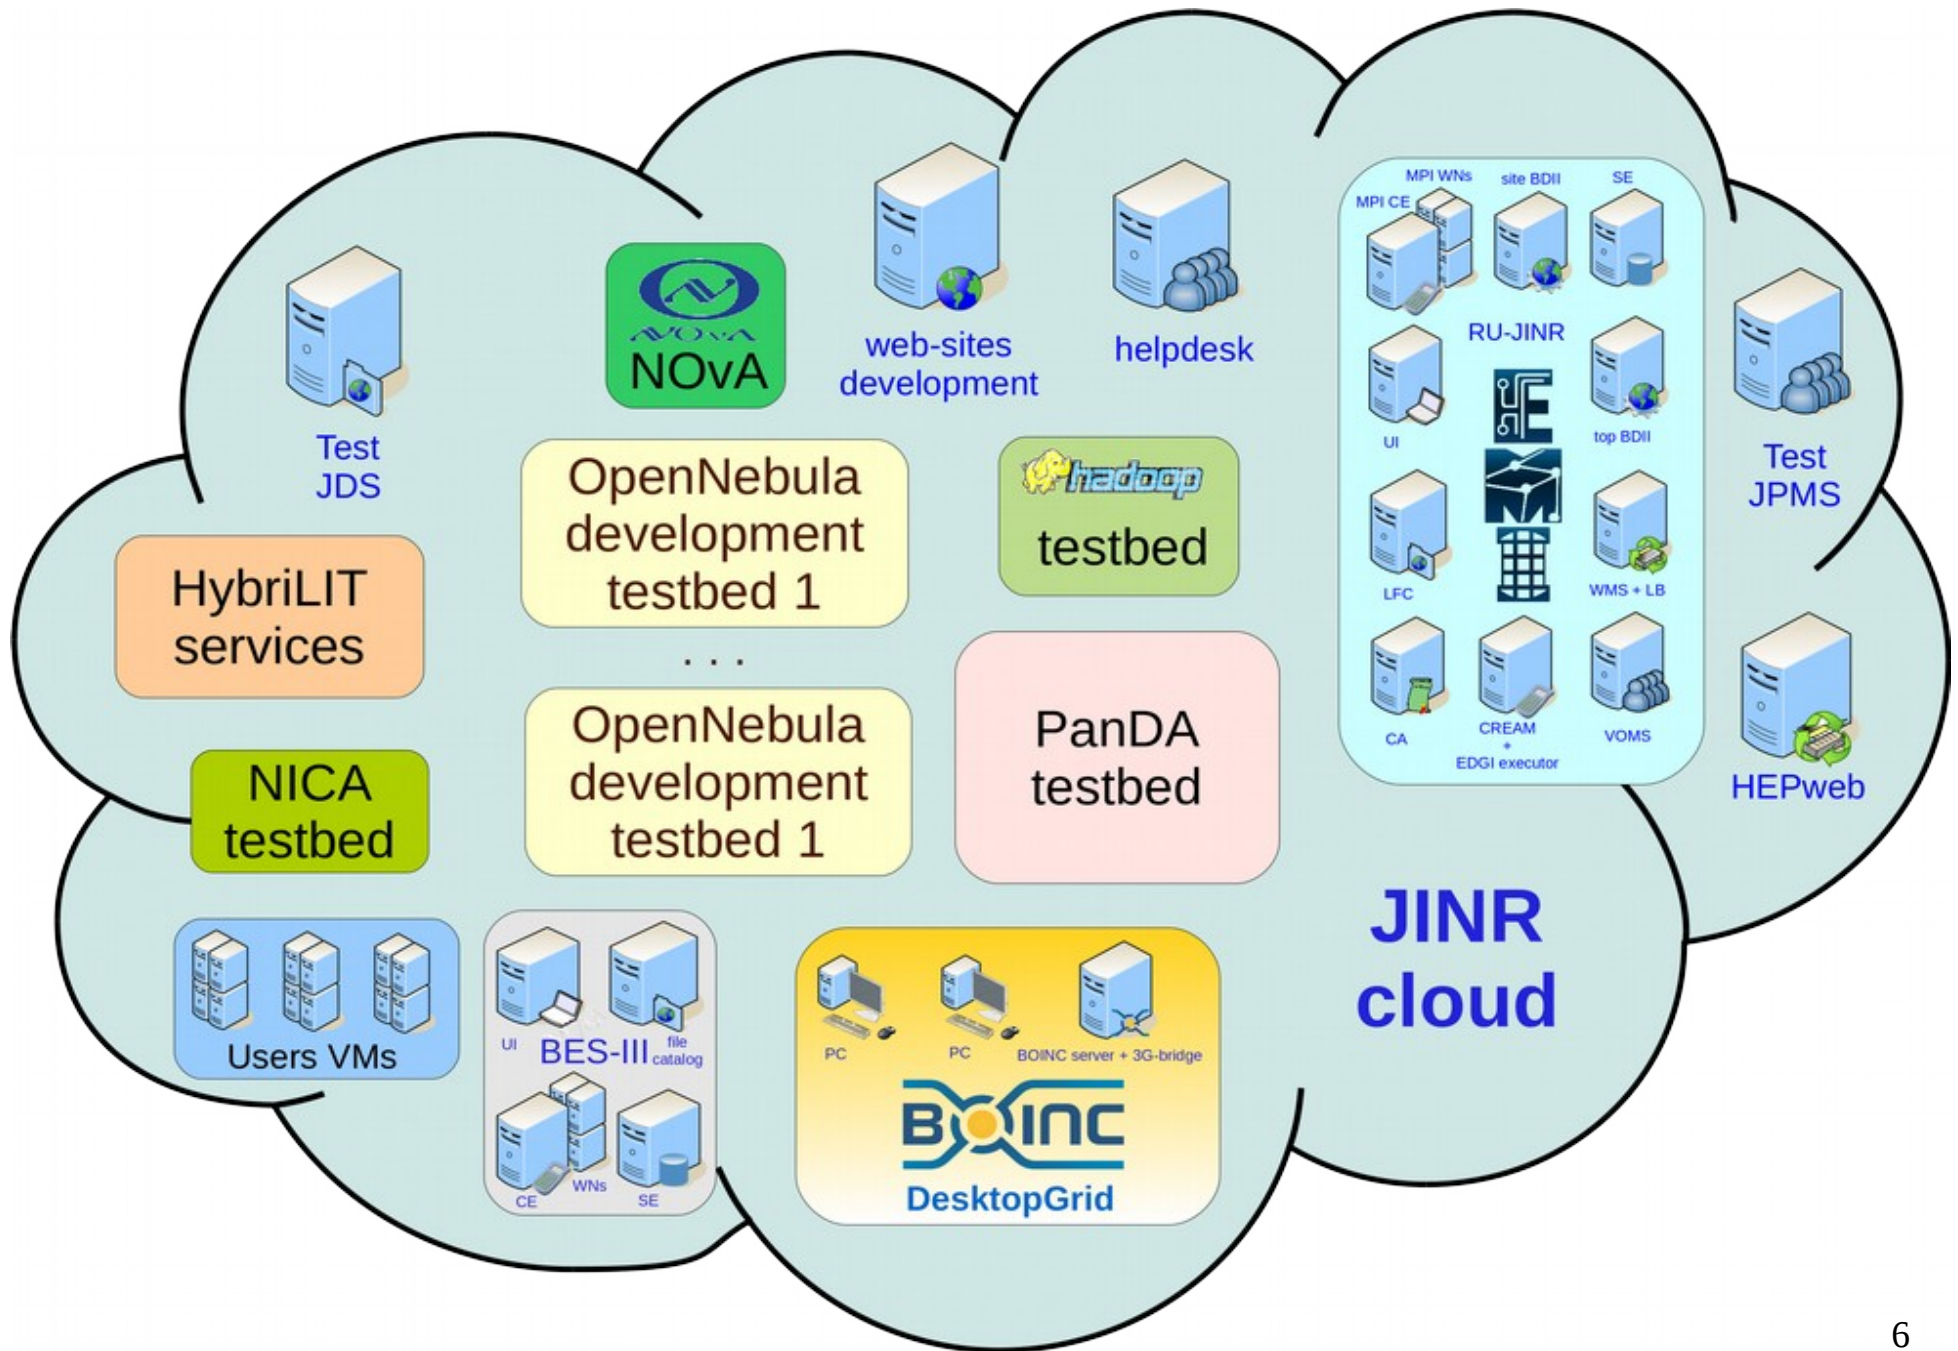
- Based on OpenNebula
- Front-end node (FN) deployed on VM:
 - opennebula core
 - MySQL DB backend
 - cloud web- and command line interfaces
 - all cloud datastores on local FS
 - web-GUI («Sunstone») with kerberos authentication + https
- 20 cluster nodes (CNS):
 - KVM VMs (6 CNS) and OpenVZ (14 CNS) containers
 - 5 OpenVZ CNS have two mirrored disks for highly reliable VMs
 - 9 OpenVZ CNS have single disk for educational, R&D or test VMs
 - VM images are on CNS' local disk copied from/to FN via ssh
- VMs
 - public IPs only
 - {rsa,dsa}-keys
 - kerberos authentication + sudo to gain root access

Service URL: <https://cloud.jinr.ru>

Characteristics

- Number of users: 74
- Number of running VMs: about 80
- Number of cores: 122
- Occupied by VMs: 134
- Available for new VM: 0
- Total RAM capacity: 252 GB
- RAM occupied by VMs: 170 GB
- RAM available for new VM: 82 GB
- Total capacity of cloud datastores: 500 GB (280 GB in use)
- User support is done via helpdesk service (helpdesk.jinr.ru)

Usage



Trainings

Organization	Organization location	Training dates	Number of trainees	Training type	Comment
JINR	Dubna, Russia	26-27.01.2015	11	usage	HybriLIT, NOvA
Gdansk university of technologies	Gdansk, Poland	06.10-12.12.2014	1	usage and administration	
National Scientific and Educational Centre of Particle and High Energy Physics of the Belarusian State University	Minsk, Belarus	22-29.09.2014	3	usage and administration	

User's workflow: first steps

OpenNebula

Request resources

Send

General Information

Full name

E-mail

Manager's full name

Manager's e-mail

Laboratory: VBLHEP

Topic number: No topic

Details on the listed topics see [here](#)

Required Resources

CPUs (cores), total per all VMs

RAM (GB), total per all VMs

Storage (GB), total per all VMs

Number of virtual machines

OS type: Linux

Comment (purpose of the requested resources or reason for quotas change)

- Dashboard
- Virtual Resources
- Infrastructure
- Marketplace
- OneFlow
- Request resources
- Statistics

- User logs into Sunstone using own kerberos credentials
- Initially (s)he has zero resource quotas
- To increase ones (s)he needs to send a request
- A ticket in JINR project management service (JPMS) is created
- JINR cloud service manager either clarify some details with user or asks cloud admin to change quotas
- User is notified via email if his/her request was accepted or rejected

User's workflow: further steps (1/4)

The screenshot shows the OpenNebula web interface for creating a virtual machine. At the top left is the 'one' logo. A navigation bar contains icons for 'Create', 'VMs', 'Templates', 'telecast', 'Log out', and 'OpenNebula'. Below this is a 'Create Virtual Machine' section with a text input field for 'Virtual Machine Name'. The main area is titled 'Select Template' and features a search bar and a refresh icon. A table lists available templates with columns for 'System' and 'Saved'. The templates include 'auth_krb', 'OpenVZ - scientific_6-x86_64 - 1CPU - 512Mb RAM - 2GB HDD_clst36', 'OpenVZ - scientific_6-x86_64 - 1CPU - 1Gb RAM - 10GB HDD_clst36', 'KVM - scientific_6-x86_64 - 1CPU - 512Mb RAM_clstKVM', 'KVM - scientific_6-x86_64 - 1CPU - 1Gb RAM_clstKVM', and 'Copy of KVM - scientific_6-x86_64 - 1CPU - 1Gb RAM_clstKVM'. Each template card shows a file icon, a size (e.g., 'x1 - 512MB'), and a '...' menu icon. A pagination bar at the bottom of the table shows '« 1 2 3 » 6'. Below the table is a 'Customize' button with a pencil icon. At the bottom center is a large blue 'Create' button.

one

Create VMs Templates telecast Log out OpenNebula

Create Virtual Machine

Virtual Machine Name

Select Template Search

System	Saved
auth_krb x1 - 512MB	
OpenVZ - scientific_6-x86_64 - 1CPU - 512Mb RAM - 2GB HDD_clst36 x1 - 512MB	
OpenVZ - scientific_6-x86_64 - 1CPU - 1Gb RAM - 10GB HDD_clst36 x1 - 1GB	
KVM - scientific_6-x86_64 - 1CPU - 512Mb RAM_clstKVM x1 - 512MB	
KVM - scientific_6-x86_64 - 1CPU - 1Gb RAM_clstKVM x1 - 1GB	
Copy of KVM - scientific_6-x86_64 - 1CPU - 1Gb RAM_clstKVM x1 - 1GB	

« 1 2 3 » 6

Customize

or

Create

User's workflow: further steps (2/4)

The screenshot displays the 'Update VM Template' wizard in OpenNebula. The interface is in 'Advanced mode' and shows the following configuration details:

- Wizard Mode:** Wizard (selected), Advanced mode
- Tabs:** General (selected), Storage, Network, OS Booting, Input/Output, Context, Other
- Description:** Windows SERV
- Memory:** 512 MB
- CPU:** 1
- Advanced options:** VCPU: 1

An 'Update' button is located at the bottom right of the wizard.

User's workflow: further steps (3/4)



Create

VMs

Templates

telecast

Log out

OpenNebula

Virtual Machines



KVM_host x1 - 1GB Public_openvz_scientific_6-x86_64_cntxt_clst36 159.93.36.129 RUNNING 1 Jun	KVM_host_2 x1 - 1GB Public_kvm_SL-6_x86-64_cntxt_clstrKVM 159.93.36.174 RUNNING 1 Jun	test-795 x1 - 2GB Public_openvz_scientific_6-x86_64_cntxt_clst36 159.93.36.130 RUNNING 1 May
Git_server x0.1 - 256MB My_openvz_scientific_6-x86_64 159.93.36.117 RUNNING 5 Apr	KVM_SL_6.4_x86-64 x1 - 512MB Public_kvm_SL-6_x86-64_cntxt_clstrKVM 159.93.36.135 RUNNING 3 Apr	KVM_OpenVZ_HOST2 x2 - 2GB KVM_OpenVZ_Host_new 159.93.36.143 RUNNING 3 Feb
FN_opennebula 4.4.1->4.6.0_test_infr x1 - 1GB FN_scientific_6-x86_64_cntxt 159.93.36.32 RUNNING 1 Feb	KVM_OpenVZ_HOST1 x2 - 2GB KVM_OpenVZ_Host_new 159.93.36.158 RUNNING 5 Feb	FN_opennebula 4.4.1_test_infr x1 - 1GB FN_scientific_6-x86_64_cntxt 159.93.36.123 RUNNING 4 Feb

Create Virtual Machine

User's workflow: further steps (4/4)

The screenshot displays a VNC connection window titled "VNC Connected (unencrypted) to: QEMU (one-780)". The desktop environment is Windows 7, running in "Test Mode" (Build 7601). The Start menu is open, showing a list of applications including Notepad, Trassir 3 Server, Getting Started, Windows Media Center, Calculator, Sticky Notes, Snipping Tool, Paint, Remote Desktop Connection, and Mozilla Firefox. A command prompt window is open in the background, showing the following commands and output:

```
C:\Windows\system32\cmd.exe
C:\>cd C:\Download\usbip_self
C:\Download\usbip_self>usbip -a 159.93.40.78 3-2
Attached to usbhub port 1
windows.c: 444 (write_to_dev) read from sock r
usbip_self>pause
continue . . .
```

The taskbar at the bottom shows the Start button, Internet Explorer, File Explorer, and Firefox. The system tray on the right indicates the time is 12:21 PM on 6/20/2014.

Monitoring

Nagios®

General

Home
Documentation

Current Status

Tactical Overview
Map
Hosts
Services
Host Groups
Summary
Grid
Service Groups
Summary
Grid
Problems
Services (Unhandled)
Hosts (Unhandled)
Network Outages

Quick Search:

Reports

Availability
Trends
Alerts
History
Summary
Histogram
Notifications
Event Log

System

Comments
Downtime
Process Info
Performance Info
Scheduling Queue
Configuration

Current Network Status
Last Updated: Mon May 25 17:08:51 MSK 2015
Updated every 90 seconds
Nagios® Core™ 3.5.1 - www.nagios.org
Logged in as test

Host Status Totals
Up Down Unreachable Pending
23 0 0 0
All Problems All Types
0 23

Service Status Totals
Ok Warning Unknown Critical Pending
291 0 0 0 0
All Problems All Types
0 291

View History For all hosts
View Notifications For All Hosts
View Host Status Detail For All Hosts

Service Status Details For All Hosts

Limit Results: 100

Host	Service	Status	Last Check	Duration	Attempt	Status Information
cldwn02	CPU Count	OK	05-25-2015 17:05:34	399d 5h 7m 20s	1/4	4
	CPU Model	OK	05-25-2015 17:06:01	419d 7h 17m 58s	1/4	model name: Intel(R) Xeon(R) CPU E5-2660 @ 2.60GHz
	Current Load	OK	05-25-2015 17:06:22	419d 7h 17m 34s	1/4	OK - load average: 2.04, 1.91, 1.90
	Current Users	OK	05-25-2015 17:06:45	419d 7h 22m 10s	1/4	USERS OK - 0 users currently logged in
	Disk Partition Root	OK	05-25-2015 17:07:09	419d 7h 21m 46s	1/4	DISK OK - free space: / 12535 MB (87% inode=96%):
	Disk Partition VZ	OK	05-25-2015 17:07:34	399d 5h 7m 53s	1/4	DISK OK - free space: /vz 148427 MB (73% inode=99%):
	PING	OK	05-25-2015 17:03:56	4d 5h 54m 55s	1/4	PING OK - Packet loss = 0%, RTA = 0.31 ms
	RAID	OK	05-25-2015 17:08:20	419d 7h 20m 34s	1/4	RAID OK - raid1: 0, errors: 0
	RAM Count	OK	05-25-2015 17:08:44	399d 5h 8m 11s	1/4	1.116.5777.68 (free/used/total) GB
	SSH Server	OK	05-25-2015 17:04:09	165d 6h 55m 56s	1/4	SSH OK - OpenSSH_5.3 (protocol 2.0)
	SWAP	OK	05-25-2015 17:04:31	165d 6h 55m 34s	1/4	SWAP OK - 87% free (7079 MB out of 8191 MB)
Total Processes	OK	05-25-2015 17:04:55	165d 6h 51m 10s	1/4	PROCS OK: 141 processes with STATE = RSZDT	
Uname	OK	05-25-2015 17:05:19	399d 5h 7m 29s	1/4	Linux cldwn02.jinr.ru 2.6.32-431.el6.x86_64 #1 SMP Thu Nov 21 13:35:52 CST 2013 x86_64 x86_64 x86_64 GNU/Linux	
cldwn03	CPU Count	OK	05-25-2015 17:05:35	131d 4h 0m 24s	1/4	2
	CPU Model	OK	05-25-2015 17:06:01	131d 4h 5m 5s	1/4	model name: Intel(R) Core(TM)2 Duo CPU E8400 @ 3.00GHz
	Current Load	OK	05-25-2015 17:06:23	131d 4h 4m 37s	1/4	OK - load average: 0.05, 0.06, 0.05
	Current Users	OK	05-25-2015 17:06:46	131d 4h 1m 0s	1/4	USERS OK - 0 users currently logged in
	Disk Partition Root	OK	05-25-2015 17:07:10	131d 4h 0m 53s	1/4	DISK OK - free space: / 11717 MB (61% inode=95%):
	Disk Partition VZ	OK	05-25-2015 17:07:34	131d 4h 0m 29s	1/4	DISK OK - free space: /vz 185446 MB (95% inode=99%):
	PING	OK	05-25-2015 17:07:58	125d 7h 35m 0s	1/4	PING OK - Packet loss = 0%, RTA = 0.26 ms
	RAID	OK	05-25-2015 17:08:24	131d 4h 0m 16s	1/4	RAID OK - raid1: 0, errors: 0
	RAM Count	OK	05-25-2015 17:08:45	131d 4h 5m 20s	1/4	0.643.09/3.74 (free/used/total) GB
	SSH Server	OK	05-25-2015 17:04:09	131d 4h 4m 55s	1/4	SSH OK - OpenSSH_5.3 (protocol 2.0)
	SWAP	OK	05-25-2015 17:04:32	131d 4h 0m 33s	1/4	SWAP OK - 100% free (8184 MB out of 8191 MB)
Total Processes	OK	05-25-2015 17:04:56	131d 4h 5m 9s	1/4	PROCS OK: 108 processes with STATE = RSZDT	
Uname	OK	05-25-2015 17:05:20	131d 4h 0m 43s	1/4	Linux cldwn03.jinr.ru 2.6.32-504.1.3.el6.x86_64 #1 SMP Tue Nov 11 14:19:04 CST 2014 x86_64 x86_64 x86_64 GNU/Linux	
cldwn04	CPU Count	OK	05-25-2015 17:05:36	399d 5h 7m 20s	1/4	2
	CPU Model	OK	05-25-2015 17:06:01	399d 8h 21m 54s	1/4	model name: Intel(R) Core(TM)2 Duo CPU E8400 @ 3.00GHz
	Current Load	OK	05-25-2015 17:06:24	399d 8h 21m 32s	1/4	OK - load average: 0.05, 0.20, 0.09
	Current Users	OK	05-25-2015 17:06:50	157d 23h 39m 14s	1/4	USERS OK - 0 users currently logged in
	Disk Partition Root	OK	05-25-2015 17:07:15	157d 23h 37m 54s	1/4	DISK OK - free space: / 7425 MB (77% inode=94%):
	Disk Partition VZ	OK	05-25-2015 17:07:35	157d 23h 37m 27s	1/4	DISK OK - free space: /vz 416015 MB (96% inode=99%):

- <http://cloud-mon.jinr.ru>
- All CNs + FN
- Notifications via email and SMS mobile phone

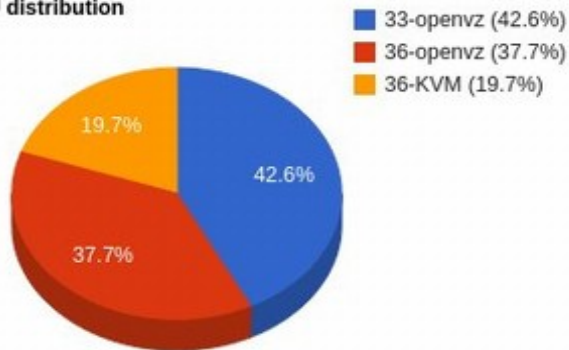
Statistics (1/2)




Statistics

Resources distribution over clusters (graphical view)

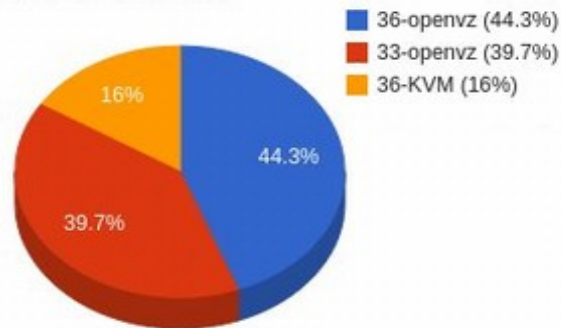
CPU distribution




Memory distribution



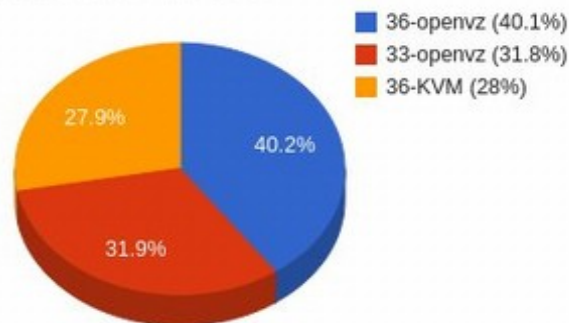
Allocated CPU distribution




Allocated Memory distribution



Distribution of CPU utilization




Distribution of Memory utilization



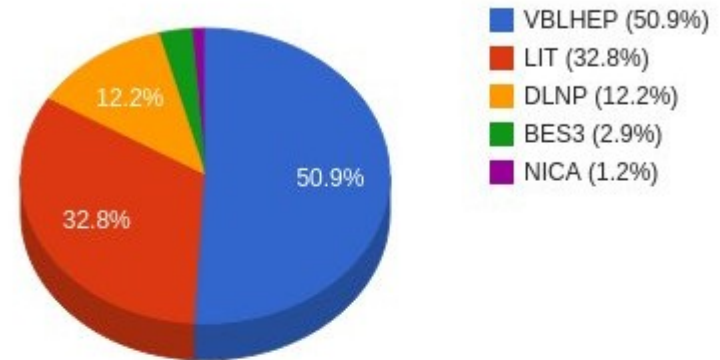
Statistics (2/2)

Resources usage by department (graphical view)

CPU usage by department, core * hours



Memory usage by department, GB * hours

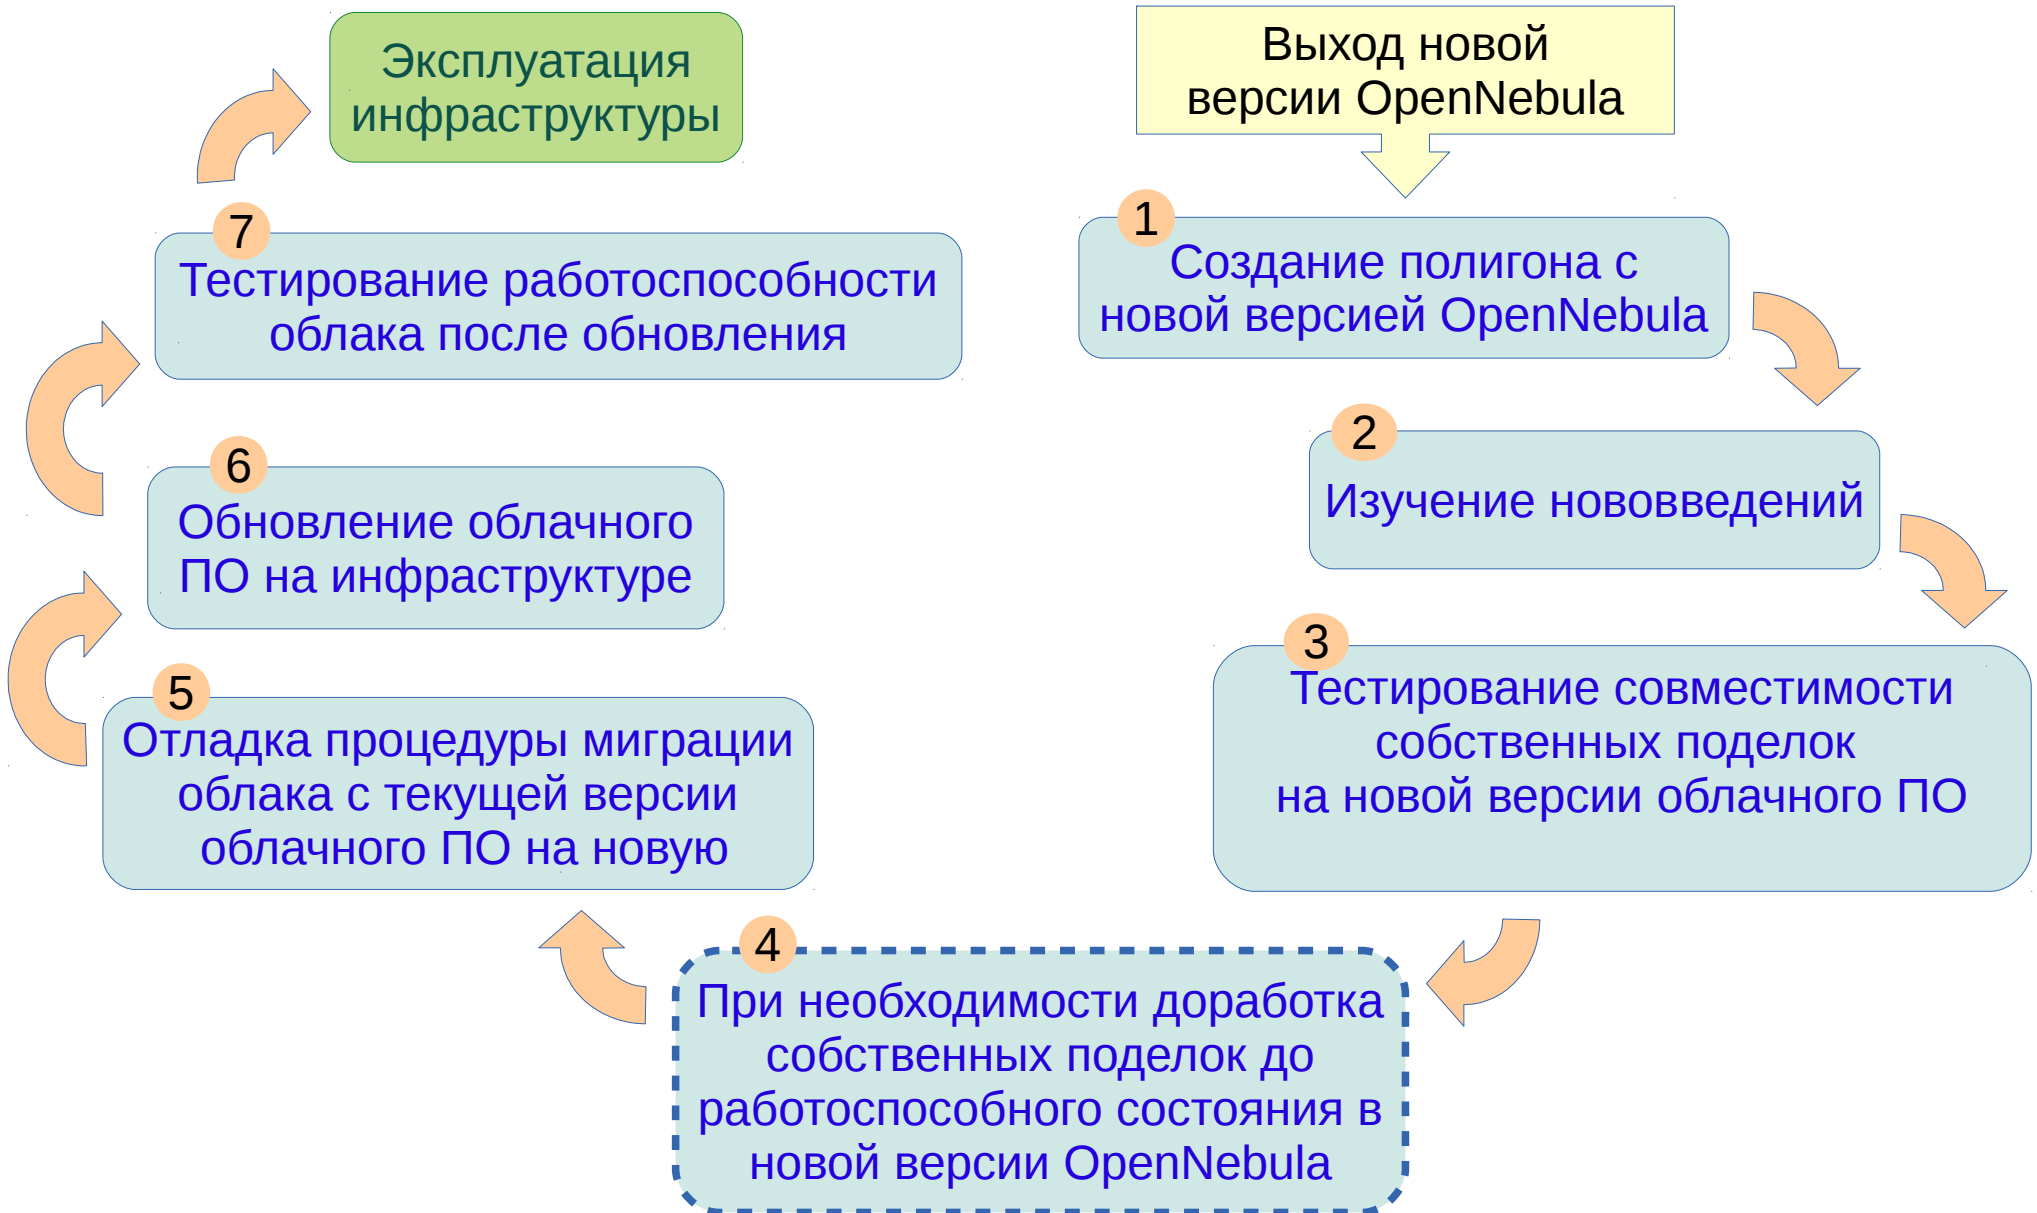


Resources usage by department

Start date: End date: [Get Accounting](#)

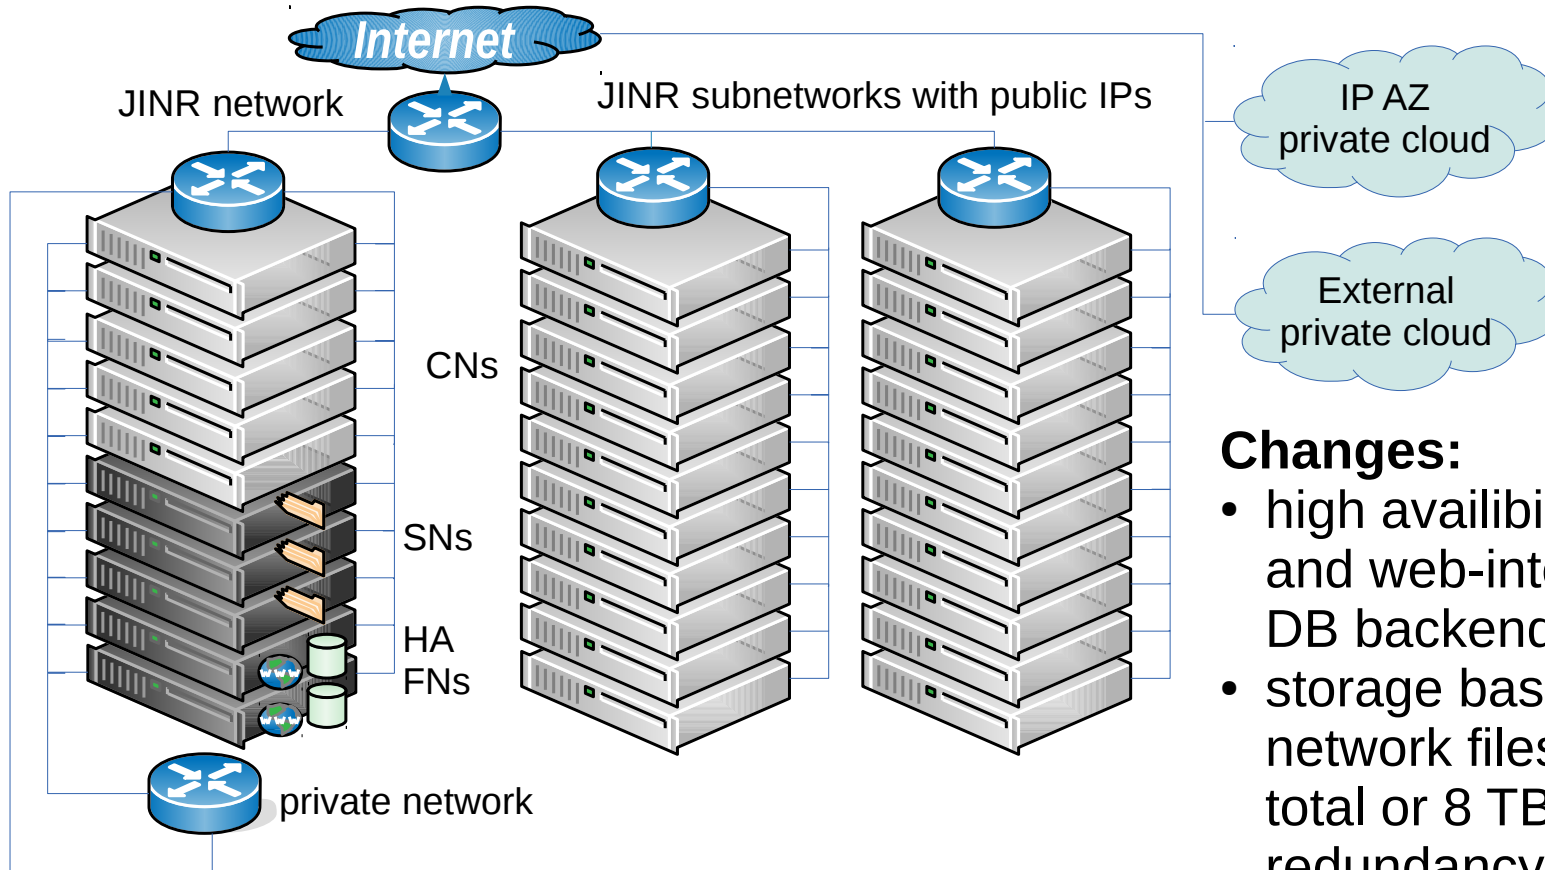
Lab Name	CPU, Core * hours	RAM, GB * hours
BES3	3534	7068
DLNP	15120	29520
LIT	69935	79493
NICA	2880	2880
TMPK	0	0
VBLHEP	120240	123120

Процесс перехода на новую версию ПО



JINR cloud service: development

It is planned to migrate existing JINR cloud service from simple configuration to highly-available one and then join resources with organisations from JINR member states.



HA FNs — highly-available front-end nodes
CNS — cloud nodes
SNs — storage nodes
IP AZ — Institute of physics (Azerbaijan)

Changes:

- high availability for cloud core and web-interfaces as well as DB backend
- storage based on distributed network filesystem (16 TB in total or 8 TB with redundancy=2)
- +80 cores, +160 GB of RAM
- +VMs with private IPs
- connected with external private clouds

Team

- Nikita Balashov
 - OpenNebula extensions development and support
 - User support and training
- Alexandr Baranov
 - Cloud administration, OpenNebula extensions testing
 - User support and training
- Nikolay Kutovskiy
 - Project coordination
 - User support
- Roman Semenov
 - Administration
 - User support